From H&M to Gap for Lightweight BWT Merging

Giovanni Manzini[1,2]

[1]Computer Science Institute, University of Eastern Piedmont, Italy
[2]IIT-CNR, Pisa, Italy

arXiv:1609.04618v1 [cs.DS] 15 Sep 2016

Abstract

Recently, Holt and McMillan [Bionformatics 2014, ACM-BCB 2014] have proposed a simple and elegant algorithm to merge the Burrows-Wheeler transforms of a family of strings. In this paper we show that the H&M algorithm can be improved so that, in addition to merging the BWTs, it can also merge the Longest Common Prefix (LCP) arrays. The new algorithm, called Gap because of how it operates, has the same asymptotic cost as the H&M algorithm and requires additional space only for storing the LCP values.

1 Introduction

Compressed indices [17] are core components of many data intensive tools, especially in bioinformatics [12, 18]. A fundamental component of many compressed indices is the Burrows Wheeler transform (BWT) of the data to be indexed, which is often complemented by the Longest Common Prefix (LCP) array and (a sampling of) the Suffix Array. Because of the sheer size of the data involved, the construction of compressed indices is a challenging problem in itself. Although the final outcome is a *compressed* index, construction algorithms can be memory intensive and the necessity of developing *lightweight*, ie space economical, algorithms was recognized since the very beginning of the field [2, 15, 16]. An alternative to space economical algorithms are external memory construction algorithms, where the challenge is to efficiently use the abundant external memory, typically by accessing data in large blocks (see [8, 9] and references therein).

Many construction algorithms for compressed indices are designed for the case the input consists in a single sequence; yet in many applications the data to be indexed consist in a collection of distinct items (documents, web pages, chromosomes, proteins, etc.). One can concatenate such items using distinct end-of-file separators and index the resulting single sequence. However, this is possible only for small collections and from the algorithmic point of view it makes no sense to "forget" that the input consists of different items: this is an additional information that algorithms should exploit to run faster.

A case of great practical interest is the one where the input is a collection of Next Generation Sequencing (NGS) reads, which typically consists in millions of sequences of lengths ranging from a few hundreds to a few thousands symbols. In this case the use of explicit distinct separators is not feasible: A few algorithms have been therefore developed specifically for this problem. The first one is the BCR algorithm [1]. For a collection of m

strings of total length n and maximum length K, BCR uses $\mathcal{O}(m \log(mK))$ bits of space and takes $\mathcal{O}(K\mathsf{sort}(m))$ time where $\mathsf{sort}(m)$ is the time to sort m integers. BCR can be modified into an external memory algorithm that uses a negligible amount of RAM and a I/O volume of $\mathcal{O}(mK^2)$ bits. Recently, BCR has been extended to compute the LCP arrays along with the BWTs in $\mathcal{O}(K(n + \mathsf{sort}(m))$ time [3].

Two other algorithms designed for NGS reads are CX1 [13] and ropeBWT [11]. The former is designed to exploit the computing power of modern GPUs, while the latter uses a dynamic data structure (a B+ tree) to maintain partial BWTs so that its complexity is $\mathcal{O}(n \log n)$ time. A comparison of the performance of these algorithms for different sets of NGS reads are reported in [11, Table 1].

Recently, Holt and McMillan [6, 7] have presented a new approach for computing the BWT of a collections of sequences based on the concept of merging: First the BWTs of the individual sequences are computed (by any single-string BWT algorithm) and then they are merged, possibly in multiple rounds as in the standard mergesort algorithm. The idea of BWT-merging is not new [5, 19] but Holt and McMillan's merging algorithm is different, and much simpler, that the previous approaches. For a constant size alphabet the algorithm in [6] merges the BWTs of two sequences t_0, t_1 in $\mathcal{O}(n \cdot \mathsf{avelcp}_{01})$ time where $n = |\mathsf{t}_0| + |\mathsf{t}_1|$ and avelcp_{01} is the average length of the common prefix between suffixes of t_0 and t_1. The average longest common prefix is $\mathcal{O}(n)$ in the worst case but $\mathcal{O}(\log n)$ for random strings and for many real word datasets [10]. The algorithm is lightweight in the sense that it uses only $\mathcal{O}(n)$ bits in addition to the space for its input and output.

In this paper we show that the H&M (Holt and McMillan) merging algorithm can be modified so that, in addition to the BWTs, it can merges the LCP arrays as well. The new algorithm, called Gap because of how it operates, has the same asymptotic cost as H&M and uses additional space only for storing its additional input and output, ie the LCP values.

2 Notation

Let $\mathsf{t}[1, n]$ denote a string of length n over a finite alphabet Σ. As is usual in the indexing literature we assume $\mathsf{t}[n]$ is a symbol not appearing elsewhere in t and lexicographically smaller than any other symbol. We write $\mathsf{t}[i, j]$ to denote the substring $\mathsf{t}[i]\mathsf{t}[i + 1] \cdots \mathsf{t}[j]$. If $j \geq n$ we assume $\mathsf{t}[i, j] = \mathsf{t}[i, n]$. If $i > j$ or $i > n$ then $\mathsf{t}[i, j]$ is the empty string. Given two strings t and s we write $\mathsf{t} \preceq \mathsf{s}$ ($\mathsf{t} \prec \mathsf{s}$) to denote that t is lexicographically (strictly) smaller than s. As usual we assume that if t is a prefix of s then $\mathsf{t} \prec \mathsf{s}$. We denote by $\mathsf{LCP}(\mathsf{t}, \mathsf{s})$ the length of the longest common prefix between t and s.

The *suffix array* $\mathsf{sa}[1, n]$ associated to t is the permutation of $[1, n]$ giving the lexicographic order of t's suffixes, that is, for $i = 1, \ldots, n - 1$, $\mathsf{t}[\mathsf{sa}[i], n] \prec \mathsf{t}[\mathsf{sa}[i + 1], n]$. The *longest common prefix* array $\mathsf{lcp}[1, n + 1]$ is defined for $i = 2, \ldots, n$ by

$$\mathsf{lcp}[i] = \mathsf{LCP}(\mathsf{t}[\mathsf{sa}[i - 1], n], \mathsf{t}[\mathsf{sa}[i], n]), \tag{1}$$

that is, the lcp array stores the length of the longest common prefix between lexicographically consecutive suffixes. In addition, for convenience of notation, we define $\mathsf{lcp}[1] = \mathsf{lcp}[n + 1] = -1$. The *Burrows-Wheeler transform* $\mathsf{bwt}[1, n]$ of t is defined by

$$\mathsf{bwt}[i] = \begin{cases} \mathsf{t}[n] & \text{if } \mathsf{sa}[i] = 1 \\ \mathsf{t}[\mathsf{sa}[i] - 1] & \text{if } \mathsf{sa}[i] > 1. \end{cases}$$

lcp	bwt	*context*
-1	b	$
0	c	ab$
2	$	abcab$
0	a	b$
1	a	bcab$
0	b	cab$
-1		

lcp	bwt	*context*
-1	c	•
0	•	aabcabc•
1	c	abc•
3	a	abcabc•
0	a	bc•
2	a	bcabc•
0	b	c•
1	b	cabc•
-1		

id	lcp_{01}	bwt_{01}	*context*
0	-1	b	$
1	0	c	•
1	0	•	aabcabc•
0	1	c	ab$
1	2	c	abc•
0	3	$	abcab$
1	5	a	abcabc•
0	0	a	b$
1	1	a	bc•
0	2	a	bcab$
1	4	a	bcabc•
1	0	b	c•
0	1	b	cab$
1	3	b	cabc•
	-1		

Figure 1: LCP array and BWT for $\mathsf{t}_0 = \mathsf{abcab\$}$ and $\mathsf{t}_1 = \mathsf{aabcabc\bullet}$, and multi-string BWT and corresponding LCP array for the same strings. Column id shows, for each entry of $\mathsf{bwt}_{01} = \mathsf{bc\bullet cc\$aaaabbb}$ whether it comes from t_0 or t_1.

In other words, $\mathsf{bwt}[1,n]$ is the permutation of t in which the position of $\mathsf{t}[j]$ coincides with the lexicographic rank of $\mathsf{t}[j+1,n]$ (or of $\mathsf{t}[1,n]$ if $j=n$) in the suffix array. In accordance with the literature we call such string the *context* of $\mathsf{t}[j]$. See Figure 1 for an example.

The longest common prefix (LCP) array, and Burrows-Wheeler transform (BWT) are fundamental components of a wide class of compressed full text indices. We are interested in the generalization of these data structures when more than one string is involved. Let $\mathsf{t}_0[1,n_0]$ and $\mathsf{t}_1[1,n_1]$ denote a pair of strings such that $\mathsf{t}_0[n_0] = \$_0$ and $\mathsf{t}_1[n_1] = \$_1$ where $\$_0 < \$_1$ are two symbols not appearing elsewhere in t_0 and t_1 and smaller than any other symbol. One can consider the concatenation $\mathsf{t}_{01}[1,n_0+n_1] = \mathsf{t}_0\mathsf{t}_1$ and define LCP array and BWT for it. However, for algorithmic reasons it is more convenient to consider the following slightly different BWT definition, first introduced in [14] and later used for example in [1, 3, 4, 7]. Let $\mathsf{sa}_{01}[1,n_0+n_1]$ denote the suffix array of the concatenation $\mathsf{t}_0\mathsf{t}_1$. The *multi-string* BWT of t_0 and t_1, denoted by $\mathsf{bwt}_{01}[1,n_0+n_1]$, is defined by

$$\mathsf{bwt}_{01}[i] = \begin{cases} \mathsf{t}_0[n_0] & \text{if } \mathsf{sa}_{01}[i] = 1 \\ \mathsf{t}_0[\mathsf{sa}_{01}[i]-1] & \text{if } 1 < \mathsf{sa}_{01}[i] \leq n_0 \\ \mathsf{t}_1[n_1] & \text{if } \mathsf{sa}_{01}[i] = n_0+1 \\ \mathsf{t}_1[\mathsf{sa}_{01}[i]-n_0-1] & \text{if } n_0+1 < \mathsf{sa}_{01}[i]. \end{cases}$$

In other words, $\mathsf{bwt}_{01}[i]$ is the symbol preceding the i-th lexicographically larger suffix, with the exception that if $\mathsf{sa}_{01}[i] = 1$ then $\mathsf{bwt}_{01}[i] = \$_0$ and if $\mathsf{sa}_{01}[i] = n_0+1$ then $\mathsf{bwt}_{01}[i] = \$_1$. Hence, $\mathsf{bwt}_{01}[i]$ always comes from the string (t_0 or t_1) that prefixes the i-th largest suffix (see again Fig. 1).

The above notion of multi-string BWT can be immediately generalized to define $\mathsf{bwt}_{01\cdots k}$ for a family of distinct strings $\mathsf{t}_0, \mathsf{t}_1, \ldots, \mathsf{t}_k$. Essentially $\mathsf{bwt}_{01\cdots k}$ is a permutation of the symbols in $\mathsf{t}_0, \mathsf{t}_1, \ldots, \mathsf{t}_k$ such that the position in $\mathsf{bwt}_{01\cdots k}$ of $\mathsf{t}_i[j]$ is given by

the lexicographic rank of its context $\mathsf{t}_i[j+1, n_i]$ (or $\mathsf{t}_i[1, n_i]$ if $j = n_i$). Note that the multi-string BWT $\mathsf{bwt}_{01\cdots k}$ defined above is related to $\mathsf{bwt}(\mathsf{t}_0\mathsf{t}_1\cdots\mathsf{t}_k)$, ie the single-string BWT of the concatenation $\mathsf{t}_0\mathsf{t}_1\cdots\mathsf{t}_k$. Indeed, since they are both defined in terms of the suffix array of $\mathsf{t}_0\mathsf{t}_1\cdots\mathsf{t}_k$, from $\mathsf{bwt}_{01\cdots k}$ we get $\mathsf{bwt}(\mathsf{t}_0\mathsf{t}_1\cdots\mathsf{t}_k)$ replacing $\$_0$ with $\$_k$, $\$_1$ with $\$_0$, $\$_2$ with $\$_1$ and so on.

Given the concatenation $\mathsf{t}_0\mathsf{t}_1$ and its suffix array $\mathsf{sa}_{01}[1, n_0 + n_1]$, we consider the corresponding LCP array $\mathsf{lcp}_{01}[1, n_0 + n_1 + 1]$ defined as in (1) (see again Fig. 1). Note that, for $i = 2, \ldots, n_0 + n_1$, $\mathsf{lcp}_{01}[i]$ gives the length of the longest common prefix between the contexts of $\mathsf{bwt}_{01}[i]$ and $\mathsf{bwt}_{01}[i-1]$. This definition can be immediately generalized to a family of k strings to define the LCP array $\mathsf{lcp}_{01\cdots k}$ associated to the multi-string BWT $\mathsf{bwt}_{01\cdots k}$.

3 The H&M algorithm revisited

In [7] Holt and McMillan introduced a simple and elegant algorithm, we call it the H&M algorithm, to merge multi-string BWTs as defined above.

Given $\mathsf{bwt}_{01\cdots k}$ and $\mathsf{bwt}_{k+1\,k+2\cdots h}$ the algorithm computes $\mathsf{bwt}_{01\cdots h}$. The computation does not explicitly need the strings $\mathsf{t}_0, \mathsf{t}_1, \ldots, \mathsf{t}_h$ but only the BWTs to be merged. For simplicity of notation we describe the algorithm assuming we are merging two single string BWTs $\mathsf{bwt}_0 = \mathsf{bwt}(\mathsf{t}_0)$ and $\mathsf{bwt}_1 = \mathsf{bwt}(\mathsf{t}_1)$; the algorithm does not change in the general case where the input are multi string BWTs. Note also that the algorithm can be easily adapted to merge more than two BWTs at the same time, that is to compute directly $\mathsf{bwt}_{01\cdots k}$ given $\mathsf{bwt}_0, \mathsf{bwt}_1, \ldots, \mathsf{bwt}_k$.

Computing bwt_{01} amounts to sorting the symbols of bwt_0 and bwt_1 according to the lexicographic order of their contexts, where the context of symbol $\mathsf{bwt}_0[i]$ (resp. $\mathsf{bwt}_1[i]$) is $\mathsf{t}_0[\mathsf{sa}_0[i], n_0]$ (resp. $\mathsf{t}_1[\mathsf{sa}_1[i], n_1]$). By construction, the symbols in bwt_0 and bwt_1 are already sorted by context, hence to compute bwt_{01} we only need to merge bwt_0 and bwt_1 without changing the relative order of the symbols within the two sequences.

The H&M algorithm works in successive phases. After the h-th phase the entries of bwt_0 and bwt_1 are sorted on the basis of the first h symbols of their context. More formally, the output of the h-th phase is a binary vector $Z^{(h)}$ containing $n_0 = |\mathsf{t}_0|$ **0**'s and $n_1 = |\mathsf{t}_1|$ **1**'s and such that the following property holds.

Property 1. For $i = 1, \ldots, n_0$, $j = 1, \ldots n_1$ the i-th **0** precedes the j-th **1** in $Z^{(h)}$ iff

$$\mathsf{t}_0[\mathsf{sa}_0[i], \mathsf{sa}_0[i] + h - 1] \preceq \mathsf{t}_1[\mathsf{sa}_1[j], \mathsf{sa}_1[i] + h - 1] \tag{2}$$

(recall that according to our notation if $\mathsf{sa}_0[i] + h - 1 > n_0$ then $\mathsf{t}_0[\mathsf{sa}_0[i], \mathsf{sa}_0[i] + h - 1]$ coincides $\mathsf{t}_0[\mathsf{sa}_0[i], n_0]$, and similarly for t_1). □

Following Property 1 we identify the i-th **0** in $Z^{(h)}$ with $\mathsf{bwt}_0[i]$ and the j-th **1** in $Z^{(h)}$ with $\mathsf{bwt}_1[j]$ so that to $Z^{(h)}$ corresponds a permutation of bwt_{01}. Property 1 is equivalent to state that we can logically partition $Z^{(h)}$ into $b(h) + 1$ blocks

$$Z^{(h)}[1, \ell_1],\ Z^{(h)}[\ell_1 + 1, \ell_2],\ \ldots,\ Z^{(h)}[\ell_{b(h)} + 1, n_0 + n_1] \tag{3}$$

such that each block corresponds to a set of bwt_{01} symbols whose contexts are prefixed by the same length-h string (the symbols with a context of length less than h are contained in singleton blocks). Within each block the symbols of bwt_0 precede those of bwt_1, and

```
1: Initialize array F[1, |Σ|]
2: k_0 ← 1; k_1 ← 1                          ▷ Init counters for bwt_0 and bwt_1
3: for k ← 1 to n_0 + n_1 do
4:     b ← Z^(h−1)[k]                        ▷ Read bit b from Z^(h−1)
5:     if b = 0 then                         ▷ Get symbol from bwt_0 or bwt_1 according to b
6:         c ← bwt_0[k_0++]
7:     else
8:         c ← bwt_1[k_1++]
9:     end if
10:    j ← F[c]++                            ▷ Get destination for b according to symbol c
11:    Z^(h)[j] ← b                          ▷ Copy bit b to Z^(h)
12: end for
```

Figure 2: Main loop of algorithm H&M for computing $Z^{(h)}$ given $Z^{(h-1)}$. Array F is initialized so that $F[c]$ contains the number of occurrences of symbols smaller than c in bwt_0 and bwt_1 plus one. Hence, the bits stored in $Z^{(h)}$ immediately after reading symbol c are stored in positions from $F[c]$ to $F[c+1]-1$ of $Z^{(h)}$.

the context of any symbol in block $Z^{(h)}[\ell_j+1, \ell_{j+1}]$ is lexicographically smaller than the context of any symbol in block $Z^{(h)}[\ell_k+1, \ell_{k+1}]$ with $k > j$.

The H&M algorithm initially sets $Z^{(0)} = \mathbf{0}^{n_0}\mathbf{1}^{n_1}$: since the context of every bwt_{01} symbol is prefixed by the same length-0 string (the empty string), there is a single block containing all bwt_{01} symbols. At phase h the algorithm computes $Z^{(h+1)}$ from $Z^{(h)}$ using the procedure in Figure 2. For completeness we report the proof of the correctness of the H&M algorithm, which is a restatement of Lemma 3.2 in [7] using our notation.

Lemma 2. *For $h = 0, 1, 2, \ldots$ the bit vector $Z^{(h)}$ satisfies Property 1.*

Proof. We prove the result by induction. For $h = 0$, $\delta = 0, 1$ $\mathsf{t}_\delta[\mathsf{sa}_\delta[i], \mathsf{sa}_\delta[i]-1]$ is the empty string so (2) is always true and Property 1 is satisfied by $Z^{(0)} = \mathbf{0}^{n_0}\mathbf{1}^{n_1}$.

To prove the "if" part, let $h > 0$ and let $1 \le v < w \le n_0 + n_1$ denote two indexes such that $Z^{(h)}[v]$ is the i-th $\mathbf{0}$ and $Z^{(h)}[w]$ is the j-th $\mathbf{1}$ in $Z^{(h)}$. We need to show that under these assumptions inequality (2) on the lexicographic order holds.

Assume first $\mathsf{t}_0[\mathsf{sa}_0[i]] \neq \mathsf{t}_1[\mathsf{sa}_1[j]]$. The hypothesis $v < w$ implies $\mathsf{t}_0[\mathsf{sa}_0[i]] < \mathsf{t}_1[\mathsf{sa}_1[j]]$ hence (2) certainly holds.

Assume now $\mathsf{t}_0[\mathsf{sa}_0[i]] = \mathsf{t}_1[\mathsf{sa}_1[j]]$. We preliminary observe that it must be $\mathsf{sa}_0[i] \neq n_0$ and $\mathsf{sa}_1[i] \neq n_1$: otherwise we would have $\mathsf{t}_0[\mathsf{sa}_0[i]] = \$_0$ or $\mathsf{t}_1[\mathsf{sa}_1[j]] = \$_1$ which is impossible since these symbols appear only once in t_0 and t_1.

Let v', w' denote respectively the value of the main loop variable k in the procedure of Figure 2 when the entries $Z^{(h)}[v]$ and $Z^{(h)}[w]$ are written (hence, during the scanning of $Z^{(h-1)}$). The hypothesis $v < w$ implies $v' < w'$. By construction $Z^{(h-1)}[v'] = \mathbf{0}$ and $Z^{(h-1)}[w'] = \mathbf{1}$. Say v' is the i'-th $\mathbf{0}$ in $Z^{(h-1)}$ and w' is the j'-th $\mathbf{1}$ in $Z^{(h-1)}$. By the inductive hypothesis on $Z^{(h-1)}$ we have

$$\mathsf{t}_0[\mathsf{sa}_0[i'], \mathsf{sa}_0[i'] + h - 2] \preceq \mathsf{t}_1[\mathsf{sa}_1[j'], \mathsf{sa}_1[j'] + h - 2], \tag{4}$$

The fundamental observation is that, being $\mathsf{sa}_0[i] \neq n_0$ and $\mathsf{sa}_1[i] \neq n_1$, it is

$$\mathsf{sa}_0[i'] = \mathsf{sa}_0[i] + 1 \qquad \text{and} \qquad \mathsf{sa}_1[j'] = \mathsf{sa}_1[j] + 1.$$

```
1: Initialize arrays F[1, |Σ|] and Block_id[1, |Σ|]
2: k_0 ← 1; k_1 ← 1                          ▷ Init counters for bwt_0 and bwt_1
3: for k ← 1 to n_0 + n_1 do
4:     if B[k] ≠ 0 and B[k] ≠ h then
5:         id ← k                             ▷ A new block of Z^(h−1) is starting
6:     end if
7:     b ← Z^(h−1)[k]                         ▷ Read bit b from Z^(h−1)
8:     if b = 0 then                          ▷ Get symbol from bwt_0 or bwt_1 according to b
9:         c ← bwt_0[k_0++]
10:    else
11:        c ← bwt_1[k_1++]
12:    end if
13:    j ← F[c]++                             ▷ Get destination for b according to symbol c
14:    Z^(h)[j] ← b                           ▷ Copy bit b to Z^(h)
15:    if Block_id[c] ≠ id then
16:        Block_id[c] ← id                   ▷ Update block id for symbol c
17:        if B[j] = 0 then
18:            B[j] = h                       ▷ A new block of Z^(h) will start here
19:        end if
20:    end if
21: end for
```

Figure 3: Main loop of the H&M algorithm modified for the computation of the lcp values. At line 1 for each symbol c we set $\mathsf{Block_id}[c] = -1$ and $F[c]$ as in Figure 2. At the beginning of the algorithm we initialize the array $B[0, n_0 + n_1]$ as $B = 1\ 0^{n_0+n_1-1}\ 1$.

Since

$$\mathsf{t}_0[\mathsf{sa}_0[i], \mathsf{sa}_0[i] + h - 1] = \mathsf{t}_0[\mathsf{sa}_0[i]]\mathsf{t}_0[\,\mathsf{sa}_0[i'], \mathsf{sa}_0[i'] + h - 2] \quad (5)$$

$$\mathsf{t}_1[\mathsf{sa}_1[j], \mathsf{sa}_1[j] + h - 1] = \mathsf{t}_1[\mathsf{sa}_1[j]]\ \mathsf{t}_1[\mathsf{sa}_1[j'], \mathsf{sa}_1[j'] + h - 2] \quad (6)$$

combining $\mathsf{t}_0[\mathsf{sa}_0[i]] = \mathsf{t}_1[\mathsf{sa}_1[j]]$ with (4) gives us (2).

For the "only if" part assume (2) holds. We need to prove that in $Z^{(h)}$ the i-th **0** precedes the j-th **1**. If $\mathsf{t}_0[\mathsf{sa}_0[i]] < \mathsf{t}_1[\mathsf{sa}_1[j]]$ the proof is immediate. If $\mathsf{t}_0[\mathsf{sa}_0[i]] = \mathsf{t}_1[\mathsf{sa}_1[j]]$, we must have

$$\mathsf{t}_0[\mathsf{sa}_0[i] + 1, \mathsf{sa}_0[i] + h - 1] \preceq \mathsf{t}_1[\mathsf{sa}_1[j] + 1, \mathsf{sa}_1[j] + h - 1].$$

By induction, if $\mathsf{sa}_0[i'] = \mathsf{sa}_0[i] + 1$ and $\mathsf{sa}_1[j'] = \mathsf{sa}_1[j] + 1$ in $Z^{(h-1)}$ the i'-th **0** precedes the j'-th **1**. During phase h, the i-th **0** in $Z^{(h)}$ is written when processing the i'-th **0** of $Z^{(h-1)}$, and the j-th **1** in $Z^{(h)}$ is written when processing the j'-th **1** of $Z^{(h-1)}$. Since in $Z^{(h-1)}$ the i'-th **0** precedes the j'-th **1** and

$$\mathsf{bwt}_0[i'] = \mathsf{t}_0[\mathsf{sa}_0[i]] = \mathsf{t}_1[\mathsf{sa}_1[j]] = \mathsf{bwt}_1[j']$$

in $Z^{(h)}$ their relative order does not change and the i-th **0** precedes the j-th **1** as claimed. □

We now show that with a simple modification to the H&M algorithm it is possible to compute, in addition to bwt_{01}, also the LCP array lcp_{01} defined in Section 2. Our strategy

for computing LCP values consists in keeping explicit track of the logical blocks we have defined for $Z^{(h)}$ and represented in (3). More precisely, we maintain an integer array $B[1, n_0 + n_1 + 1]$ such that at the end of phase h it is $B[i] \neq 0$ iff a block of $Z^{(h)}$ starts at position i. The use of such integer array is shown in Figure 3. Note that: (i) initially we set $B = 1\,0^{n_0+n_1-1}\,1$ and once an entry in B becomes nonzero it is never changed, (ii) during phase h we only write to B the value h, (iii) in the test at Line 4 the value h is equivalent to 0, hence the values written during phase h influence the algorithm only in subsequent phases. The following lemma shows that the nonzero values of B at the end of phase h mark the boundaries of $Z^{(h)}$'s logical blocks.

Lemma 3. *For any $h \geq 0$, let ℓ, m be such that $1 \leq \ell \leq m \leq n_0 + n_1$ and*

$$\mathsf{lcp}_{01}[\ell] < h, \quad \min(\mathsf{lcp}_{01}[\ell+1], \ldots, \mathsf{lcp}_{01}[m]) \geq h, \quad \mathsf{lcp}_{01}[m+1] < h. \tag{7}$$

Then, at the end of phase h the array B is such that

$$B[\ell] \neq 0, \quad B[\ell+1] = \cdots = B[m] = 0, \quad B[m+1] \neq 0 \tag{8}$$

and $Z^{(h)}[\ell, m]$ is one of the blocks in (3)*.*

Proof. We prove the result by induction on h. For $h = 0$, hence before the execution of the first phase, (7) is only valid for $\ell = 1$ and $m = n_0 + n_1$ (recall we defined $\mathsf{lcp}_{01}[1] = \mathsf{lcp}_{01}[n_0 + n_1 + 1] = -1$). Since initially $B = 1\,0^{n_0+n_1-1}\,1$ our claim holds.

Suppose now that (7) holds for some $h > 0$. Let $s = \mathsf{t}_{01}[\mathsf{sa}_{01}[\ell], \mathsf{sa}_{01}[\ell] + h - 1]$; by (7) s is a common prefix of the suffixes starting at positions $\mathsf{sa}_{01}[\ell]$, $\mathsf{sa}_{01}[\ell+1]$, $\ldots$, $\mathsf{sa}_{01}[m]$, and no other suffix of t_{01} is prefixed by s. By Property 1 the **0**'s and **1**'s in $Z^{(h)}[\ell, m]$ corresponds to the same set of suffixes That is, if $\ell \leq v \leq m$ and $Z^{(h)}[v]$ is the ith **0** (resp. jth **1**) of $Z^{(h)}$ then the suffix starting at $\mathsf{t}_0[\mathsf{sa}_0[i]]$ (resp. $\rightarrow [\mathsf{sa}_1[j]]$) is prefixed by s.

To prove (8) we start by showing that, if $\ell < m$, then at the end of phase $h - 1$ it is $B[\ell+1] = \cdots = B[m] = 0$. To see this observe that the range $\mathsf{sa}_{01}[\ell, m]$ is part of a (possibly) larger range $\mathsf{sa}_{01}[\ell', m']$ containing all suffixes prefixed by the length $h-1$ prefix of s. By inductive hypothesis, at the end of phase $h - 1$ it is $B[\ell'+1] = \cdots = B[m'] = 0$ which proves our claim since $\ell' \leq \ell$ and $m \leq m'$.

To complete the proof, we need to show that during phase h: (i) we do not write a nonzero value in $B[\ell+1, m]$ and (ii) we write a nonzero to $B[\ell]$ and $B[m+1]$ if they do not already contain a nonzero. Let $c = s[0]$ and $s' = s[1, h-1]$ so that $s = cs'$. Consider now the range $\mathsf{sa}_{01}[e, f]$ containing the suffixes prefixed by s'. By inductive hypothesis at the end of phase $h - 1$ it is

$$B[e] \neq 0, \quad B[e+1] = \cdots = B[f] = 0, \quad B[f+1] \neq 0. \tag{9}$$

During iteration h, the bits in $Z^{(h)}[\ell, m]$ are possibly changed only when we are scanning the region $Z^{(h-1)}[e, f]$ and we find an entry $b = Z^{(h-1)}[k]$, $e \leq k \leq f$, such that the corresponding value in bwt_b is c. Note that by (9) as soon as k reaches e the variable id changes and becomes different from all values stored in $\mathsf{Block_id}$. Hence, at the first occurrence of symbol c the value h will be stored in $B[\ell]$ (Line 18) unless a nonzero is already there. Again, because of (9), during the scanning of $Z^{(h-1)}[e, f]$ the variable id does not change so subsequent occurrences of c will not cause a nonzero value to be written to $B[\ell+1, m]$. Finally, as soon as we leave region $Z^{(h-1)}[e, f]$ and k reaches $f + 1$, the variable id changes again and at the next occurrence of c a nonzero value will be stored in

$B[m+1]$. If there are no more occurrences of c after we leave region $Z^{(h-1)}[e,f]$ then either $\mathsf{sa}_{01}[m+1]$ is the first suffix array entry prefixed by symbol $c+1$ or $m+1 = n_0+n_1+1$. In the former case $B[m+1]$ gets a nonzero value at phase 1, in the latter case $B[m+1]$ gets a nonzero when we initialize array B.

This completes the proof □

Corollary 4. *For $i = 2,\ldots,n_0+n_1$, if $\mathsf{lcp}_{01}[i] = \ell$, then starting from the end of phase $\ell+1$ it is $B[i] = \ell+1$.*

Proof. By Lemma 3 we know that $B[i]$ becomes nonzero only after phase $\ell+1$. Since at the end of phase ℓ it is still $B[i] = 0$ during phase $\ell+1$ $B[i]$ gets the value $\ell+1$ which is never changed in successive phases. □

The above corollary suggests the following algorithm to compute bwt_{01} and lcp_{01}: repeat the procedure of Figure 3 until the phase h in which all entries in B become nonzero. At that point $Z^{(h)}$ describes how bwt_0 and bwt_1 should be merged to get bwt_{01} and for $i = 2,\ldots,n_0+n_1$ $\mathsf{lcp}_{01}[i] = B[i]-1$. The above strategy requires a number of iterations, each one taking $\mathcal{O}(n_0+n_1)$ time, equal to the maximum of the lcp values for an overall complexity of $\mathcal{O}((n_0+n_1)\mathsf{maxlcp}_{01})$, where $\mathsf{maxlcp}_{01} = \max_i \mathsf{lcp}_{01}[i]$. In the next section we describe a much faster algorithm that goes beyond this simple strategy and avoids to re-process the portions of B and $Z^{(h)}$ which are no longer relevant for the computation of the final result.

4 The Gap algorithm

Definition 5. If $B[\ell] \neq 0$, $B[m+1] \neq 0$ and $B[\ell+1] = \cdots = B[m] = 0$, we say that block $Z^{(h)}[\ell,m]$ is *monochrome* if it contains only **0**'s or only **1**'s. □

Since a monochrome block only contains suffixes from either t_0 or t_1, whose relative order and LCP's are known, it does not need to be further modified. This intuition is formalized by the following lemmas.

Lemma 6. *If at the end of phase h bit vector $Z^{(h)}$ contains only monochrome blocks we can compute bwt_{01} and lcp_{01} in $\mathcal{O}(n_0+n_1)$ time.*

Proof. By Property 1, if we identify the i-th **0** in $Z^{(h)}$ with $\mathsf{bwt}_0[i]$ and the j-th **1** with $\mathsf{bwt}_1[j]$ the only elements which could be out of order (ie not correctly sorted by context) are those within the same block. However, if the blocks are monochrome all elements belongs to either bwt_0 or bwt_1 so their relative order is correct.

To compute lcp_{01} we observe that if $B[i] \neq 0$ then by (the proof of) Corollary 4 it is $\mathsf{lcp}_{01}[i] = B[i]-1$. If instead $B[i] = 0$ we are inside a block hence $\mathsf{sa}_{01}[i-1]$ and $\mathsf{sa}_{01}[i-1]$ belongs to the same string t_0 or t_1 and their lcp is directly available in lcp_0 or lcp_1. □

Lemma 7. *Suppose that, at the end of phase h, $Z^{(h)}[\ell,m]$ is a monochrome block. Then (i) for $g > h$, $Z^{(g)}[\ell,m] = Z^{(h)}[\ell,m]$, and (ii) processing $Z^{(h)}[\ell,m]$ during phase $h+1$ creates a set of monochrome blocks in $Z^{(h+1)}$.*

Proof. The first part of the Lemma follows from the observation that subsequent phases of the algorithm will only reorder the values within a block (and possibly create new sub-blocks); but if a block is monochrome the reordering will not change its actual content.

```
 1: if (next block is irrelevant) then
 2:     skip it
 3: else
 4:     process block
 5:     if (processed block is monochrome) then
 6:         mark it irrelevant
 7:     end if
 8: end if
 9: if (last two blocks are irrelevant) then
10:     merge them
11: end if
```

Figure 4: Main loop of the Gap algorithm. The processing of active blocks at Line 4 is done as in Lines 7–20 of Figure 3.

For the second part, we observe that during phase $h+1$ as k goes from ℓ to m the algorithm writes to $Z^{(h+1)}$ the same value which is in $Z^{(h)}[\ell, m]$. Hence, a new monochrome block will be created for each distinct symbol encountered (in bwt_0 or bwt_1) as k goes through the range $[\ell, m]$. □

The lemma implies that, if block $Z^{(h)}[\ell, m]$ is monochrome at the end of phase h, starting from phase $g = h+2$ processing the range $[\ell, m]$ will not change $Z^{(g)}$ with respect to $Z^{(g-1)}$. Indeed, by the lemma the monochrome blocks created in phase $h+1$ do not change in subsequent phases (in a subsequent phase a monochrome block can be split in sub-blocks, but the actual content of the bit vector does not change). The above observation suggests that, after we have processed block $Z^{(h+1)}[\ell, m]$ in phase $h+1$, we can mark it as *irrelevant* and avoid to process it again. As the computation goes on, more and more blocks become irrelevant. Hence, in the generic phase h instead of processing the whole $Z^{(h-1)}$ we process only the blocks which are still "active" and skip irrelevant blocks. Adjacent irrelevant blocks are merged so that among two active blocks there is at most one irrelevant block (the *gap* that gives the name to the algorithm). The overall structure of a single phase is shown in Figure 4. The algorithm terminates when there are no more active blocks since this implies that all blocks have become monochrome and by Lemma 6 we are able to compute bwt_{01} and lcp_{01}.

We point out that at Line 2 of the Gap algorithm we cannot simply skip an irrelevant block ignoring its content. To keep the algorithm consistent we must correctly update the global variables of the main loop, i.e. the array F and the pointers k_0 and k_1 in Figure 3. To this end a simple approach is to store for each (merged) irrelevant block the number of occurrences o_c of each symbol $c \in \Sigma$ in it and the pair (r_0, r_1) providing the number of **0**'s and **1**'s in the block. When the algorithm reaches an irrelevant block, F, k_0, k_1 are updated setting $k_0 \leftarrow k_0 + r_0$, $k_1 \leftarrow k_1 + r_1$ and $\forall c\ F[c] \leftarrow F[c] + o_c$.

The above scheme for handling irrelevant blocks is simple and probably effective in most cases. However, using $\mathcal{O}(|\Sigma|)$ time to skip an irrelevant block is not competitive in terms of worst case complexity. A better alternative is to build a wavelet tree for bwt_0 and bwt_1 at the beginning of the algorithm. Then, for each irrelevant block we store only the the pair (r_0, r_1). When we reach an irrelevant block we use such pair to update k_0 and k_1. The array F is not immediately updated: Instead we maintain two global arrays $L_0[1, |\Sigma|]$ and $L_1[1, |\Sigma|]$ such that $L_0[c]$ and $L_1[c]$ store the value of k_0 and k_1 at the time

the value $F[c]$ was last updated. At the *first* occurrence of a symbol c inside an active block we update $F[c]$ adding to it the number of occurrences of c in $\mathsf{bwt}_0[L_o[c]+1, k_0]$ and $\mathsf{bwt}_1[L_1[c]+1, k_1]$ that we compute in $\mathcal{O}(\log|\Sigma|)$ time using the wavelet trees. Using this lazy update mechanism, handling irrelevant blocks adds a $\mathcal{O}(\min(\ell, |\Sigma|)\log|\Sigma|)$ additive slowdown to the cost of processing an active block of length ℓ.

Theorem 8. *Given* $\mathsf{bwt}_0, \mathsf{lcp}_0$ *and* $\mathsf{bwt}_1, \mathsf{lcp}_1$ *the* Gap *algorithm computes* bwt_{01} *and* lcp_{01} *in* $\mathcal{O}(\log(|\Sigma|)(n_0+n_1)\mathsf{avelcp}_{01})$ *time, where* $\mathsf{avelcp}_{01} = (\sum_i \mathsf{lcp}_{01}[i])/(n_0+n_1)$ *is the average LCP of the string* t_{01}.

Proof. The correctness follows from the above discussion. For the analysis of the running time we reason as in [6] and observe that the sum, over all phases, of the length of all active blocks is bounded by $\mathcal{O}(\sum_i \mathsf{lcp}_{01}[i]) = \mathcal{O}((n_0+n_1)\mathsf{avelcp}_{01})$. In any phase, using the lazy update mechanism, the cost of processing an active block of length ℓ is bounded by $\mathcal{O}(\ell\log(|\Sigma|)$ and final time bound follows. □

We point out that our Gap algorithm is related to the version of the H&M algorithm described in [6, Sect. 2.1]: Indeed, the sorting operations are essentially the same in the two algorithms. The main difference is that Gap keeps explicit track of the irrelevant blocks while H&M keeps explicit track of the active blocks (called buckets in [6]): this difference makes the non-sorting operations completely different. An advantage of working with irrelevant blocks is that they can be easily merged, while this is not the case for the active blocks in H&M. Of course, the main difference is that Gap computes simultaneously bwt_{01} and lcp_{01} while H&M only computes bwt_{01}.

References

[1] Bauer, M.J., Cox, A.J., Rosone, G.: Lightweight algorithms for constructing and inverting the BWT of string collections. *Theor. Comput. Sci.* **483** (2013) 134–148

[2] Burkhardt, S., Kärkkäinen, J.: Fast lightweight suffix array construction and checking. In: *Proc. 14th Symposium on Combinatorial Pattern Matching (CPM '03)*, Springer-Verlag LNCS n. 2676 (2003) 55–69

[3] Cox, A.J., Garofalo, F., Rosone, G., Sciortino, M.: Lightweight LCP construction for very large collections of strings. *J. Discrete Algorithms* **37** (2016) 17–33

[4] da Louza, F.A., Telles, G.P., de Aguiar Ciferri, C.D.: External memory generalized suffix and LCP arrays construction. In: *CPM*. Volume 7922 of *Lecture Notes in Computer Science.*, Springer (2013) 201–210

[5] Ferragina, P., Gagie, T., Manzini, G.: Lightweight data indexing and compression in external memory. In: *Proc. 9th Latin American Theoretical Informatics Symposium (LATIN '10)*. Lecture Notes in Computer Science vol. 6034 (2010) 698–711

[6] Holt, J., McMillan, L.: Constructing Burrows-Wheeler transforms of large string collections via merging. In: *BCB*, ACM (2014) 464–471

[7] Holt, J., McMillan, L.: Merging of multi-string BWTs with applications. *Bioinformatics* **30** (2014) 3524–3531

[8] Kärkkäinen, J., Kempa, D.: Engineering a lightweight external memory suffix array construction algorithm. In: *ICABD*. Volume 1146 of *CEUR Workshop Proceedings.*, CEUR-WS.org (2014) 53–60

[9] Kärkkäinen, J., Kempa, D.: LCP array construction in external memory. *ACM Journal of Experimental Algorithmics* **21** (2016) 1.7:1–1.7:22

[10] Léonard, M., Mouchard, L., Salson, M.: On the number of elements to reorder when updating a suffix array. *J. Discrete Algorithms* **11** (2012) 87–99

[11] Li, H.: Fast construction of FM-index for long sequence reads. *Bioinformatics* **30** (2014) 3274–3275

[12] Li, H., Homer, N.: A survey of sequence alignment algorithms for next-generation sequencing. *Briefings in Bioinformatics* **11** (2010) 473–483

[13] Liu, C., Luo, R., Lam, T.W.: GPU-accelerated BWT construction for large collection of short reads. *CoRR* **abs/1401.7457** (2014)

[14] Mantaci, S., Restivo, A., Rosone, G., Sciortino, M.: An extension of the burrows-wheeler transform. *Theor. Comput. Sci.* **387** (2007) 298–312

[15] Manzini, G.: Two space saving tricks for linear time LCP computation. In: *Proc. of 9th Scandinavian Workshop on Algorithm Theory (SWAT '04)*, Springer-Verlag LNCS n. 3111 (2004) 372–383

[16] Manzini, G., Ferragina, P.: Engineering a lightweight suffix array construction algorithm. In: *Proc. 10th European Symposium on Algorithms (ESA)*, Springer Verlag LNCS n. 2461 (2002) 698–710

[17] Navarro, G., Mäkinen, V.: Compressed full-text indexes. *ACM Computing Surveys* **39** (2007)

[18] Reinert, K., Langmead, B., Weese, D., Evers, D.J.: Alignment of Next-Generation Sequencing Reads. *Annual review of genomics and human genetics* **16** (2015) 133–51

[19] Sirén, J.: Compressed suffix arrays for massive data. In: *Proc. 16th Int. Symp. on String Processing and Information Retrieval (SPIRE '09)*, Springer Verlag LNCS n. 5721 (2009) 63–74